

THE Tech Guide 2026

Charting the path to
the intelligent enterprise

Table of contents

2

FOREWORD
Choose your path
by Stuart Brown

3

KNOW THE DIVIDE
What's happening now

INSIGHT | 4
What counts as AI?

INSIGHT | 5
The digital brain

6

MAKE THE MOVES
A five-part model

INSIGHT | 7
AI-native process
reengineering

PODCAST | 8
Moonshot first, then tools

INSIGHT | 9
Tokens: The economics
of intelligence

PODCAST | 10
A masterclass in people-first AI

FROM THE AI STUDIO | 11
Context and ontology

INSIGHT | 12
Metrics that matter

VISUALIZATION | 13
The path to the
intelligent enterprise

14

START HERE
Enterprise intelligence
in action

Energy and utilities | 15
Financial services | 15
Government | 16
Healthcare | 16

17

CLEAR THE WAY
Misconceptions that
can stall progress

6 myths that can
slow the shift | 18
Closing thoughts | 19

20

ABOUT THE GUIDE
Methodology and
contributors

Turn the Guide
into results | 21

FOREWORD

Choose your path



If intelligence could exist everywhere in my enterprise, how would we operate differently?

Right now, a leader in your industry is making a decision about AI that will matter more than any technology they buy this decade.

They're choosing a path.

One path applies AI tools on top of the organization's existing work. The gains are real, and they arrive quickly. Reports are drafted in minutes. Service tickets clear sooner. Teams save hours every week. These gains matter, and they're worth capturing. They're also incremental by design because the work's structure stays the same.

The other path treats AI as a reason to redesign the work itself. It starts by asking a different question: If intelligence could exist everywhere in my enterprise, how would we operate differently? On this path, leaders rethink how systems function and rebuild processes around desired outcomes. This is the intelligent

enterprise: an organization whose knowledge, judgment, and execution are connected, with expertise available wherever work happens.

Both paths are open today. Deliberately choosing between them is the defining leadership decision of this moment.

We built this year's Tech Guide to help you with that choice and what comes after. It draws on the research, applied knowledge, and on-the-ground experience of Guidehouse technology experts working across business and government organizations.

We start with the split itself and where the opportunity has moved: from digitizing information to digitizing decisions and actions. From there, we lay out our model: five checkpoints for operationalizing intelligence in complex, real-world workflows—introduced in sequence but pursued in tandem.

We then bring the model to life with a look at where intelligence is creating value across industries along with practical use cases for a deeper dive.

And finally, we tackle misconceptions that can stall momentum, from treating AI as a technology problem to mistaking a portfolio of pilots for a strategy.

AI is moving faster than any shift most leaders have managed before. In complex, regulated industries, speed matters, but speed without assurance creates new risk. The organizations that lead will treat AI not as a tool to install but as a transformation of how work is designed, governed, and delivered. That's the path to the intelligent enterprise.

Let's get started.



Stuart Brown
Chief Transformation Officer
Guidehouse

Stuart Brown, Guidehouse Chief Transformation Officer (CXO), is a celebrated technologist and thought leader with nearly three decades of transformative leadership experience across diverse industries. As CXO, Stuart leads the firm's transformation strategy, connecting industry expertise, frontier technology, and new delivery models to help clients modernize operations and realize value faster.

1 Know the divide

What's happening now

The AI conversation has entered a new phase

for the past two years, organizations have focused on AI adoption. They've deployed pilots, experimented with generative AI, and introduced new tools designed to help employees work faster. Those investments have produced meaningful gains in productivity and efficiency. But a divide is emerging. Many organizations are applying AI to existing work. Others are using it to rethink how work happens in the first place. At the same time, a broader shift is underway. Organizations are moving beyond digitizing information and toward digitizing decisions, actions, and execution. Understanding these two trends is essential to seeing where value will come from next.

The two paths

Organizations are pursuing two different paths with AI.

The first focuses on improving existing work. Employees use AI to summarize information, accelerate research, generate content, and automate routine work in ways that were impossible only a few years ago. These improvements create real value and help organizations build AI skills across the workforce.

The second path focuses on redesigning how work flows through an organization. Processes and operating models are rebuilt with AI in mind from the start. Intelligence is embedded into the day-to-day execution of work, helping organizations move knowledge more effectively, unlock expertise more broadly, and deliver outcomes with greater consistency and scale.

Think of the difference this way: Adding GPS to a taxi helps the driver find a faster route. The destination, the business model, and the operating process remain largely unchanged. This is productivity.

But building a ride-sharing platform? That transforms how transportation is coordinated and delivered. It redesigns the workflow by connecting riders, drivers, data, and decision-making into a single digital system. Autonomous vehicles aim to extend that evolution even further, embedding intelligence directly into execution. That's transformation.

The real divide isn't between AI adopters and laggards. Nearly every organization has adopted something. The difference is between deploying tools onto the existing operating model and redesigning the operating model around intelligent execution. In high-stakes, regulated industries, that difference becomes even more consequential. Generic AI layered onto existing work can't substitute for embedded, domain-specific intelligence with clear accountability.

This is assured transformation: redesigning work and embedding intelligence in ways that move at speed without outrunning security, governance, workforce readiness, resilience, or mission accountability.



INSIGHT What counts as AI?

AI is often discussed as if it were a single technology. In practice, it spans a broad spectrum of capabilities. Generative AI may dominate headlines, yet organizations are creating value through machine learning, predictive models, intelligent automation, and emerging agentic systems. It's best to think of AI as a class of machine-based systems that produce inferences from data and generate outputs such as summaries or recommendations.

What distinguishes AI from traditional software is how it adapts, recognizes patterns, and supports complex decision-making under changing conditions. The common thread is its ability to transform information into action.



Enter the era of digital execution

This divide reflects a broader shift in the role of technology within the enterprise.

For more than a decade, digital transformation focused on digitizing information. Organizations migrated workloads to the cloud, integrated applications, and increased data accessibility. These investments created the digital foundation required for the next wave of change.

Today, the focus is shifting from digitizing information to digitizing *execution*. AI now shapes how decisions are made, how work is coordinated, and how knowledge moves across teams and systems. Increasingly, intelligence is becoming embedded directly into workflows instead of sitting alongside them.

A new layer of organizational intelligence is emerging, often called a “digital brain” (see *Insight sidebar*). Knowledge that once lived in documents, systems, and individual experts can now be captured, connected, and applied across workflows. Leading organizations are treating that operating know-how as a strategic asset rather than an implementation detail.

Many organizations are applying AI to existing work. Others are using it to rethink how work happens in the first place.

 INSIGHT
The digital brain

Every organization has a form of institutional intelligence: the collective knowledge contained in its data, decisions, processes, policies, and people. Yet much of that intelligence remains fragmented across systems and difficult to access.

The concept of a digital brain describes an organization’s ability to capture, connect, and operationalize that intelligence at scale. Instead of treating data as a byproduct of operations, leading organizations are increasingly treating it as strategic intellectual property. The goal is to make knowledge easier to access and apply across work so teams can make faster decisions, reduce duplicated effort, and preserve critical expertise when employees leave.

Building a digital brain begins with creating a shared layer of organizational knowledge. Common approaches include establishing an enterprise ontology (a shared knowledge base), using vector databases and knowledge graphs to organize information, and implementing retrieval-augmented generation (RAG) that grounds AI reasoning in trusted internal content.

Together, these components turn the digital brain into an institutional asset: a governed combination of enterprise knowledge, operating context, decision rights, and accumulated judgment. The organizations creating the greatest value from AI are going beyond deploying models. They’re building systems that allow knowledge to flow through the enterprise and putting stewards in charge of what those systems know.



2 Make the moves

A five-part model

Once you choose a path, checkpoints help you stay on course

to bring the intelligent enterprise to life and build the digital brain, you must choose the path of enterprise redesign and commit to it deliberately. Like any journey, the path includes a series of checkpoints: places to pause, confirm the direction, and keep going. The five that follow are the common threads found in organizations successfully operationalizing intelligence. They're presented in sequence, but they advance in parallel. Together, they provide a practical path for operationalizing intelligence in complex workflows, creating value that compounds over time.



CHECKPOINT 1

Define the work



The most successful organizations start smaller and think bigger.

Start with an outcome, then redesign around it

Organizations often begin their AI journey with a list of possibilities. The challenge is that value becomes difficult to measure when attention is spread across dozens of ideas at once.

The most successful organizations start smaller and think bigger.

Rather than launching a portfolio of pilots, identify a single mission-critical workflow and aim to redesign it from the ground up. The workflow should matter to the business, touch multiple systems or teams, and produce a measurable outcome. Examples include customer onboarding, ambient listening in a healthcare setting, processing mortgage applications, or managing outage response for a utility provider.

This checkpoint is about defining success *upfront*—before selecting tools. Establish what you want to achieve, how the workflow performs today, and how you'll measure success. Speed, quality, customer experience, cost, and mission impact are all common starting points. Organizations that front-load these discussions create a stronger foundation for every decision that follows.

The workflow itself also becomes a useful lens for reengineering (*see Insight sidebar*). Most business processes extend across multiple systems, handoffs, and decisions. They often contain hidden delays, duplicated effort, and knowledge trapped in silos. Focusing on a single workflow makes those friction points visible and creates opportunities to rethink how the work gets done.



INSIGHT AI-native process reengineering

Once a mission-critical workflow has been identified, many leading organizations take the additional step of process reengineering. This approach begins with a thought exercise: temporarily pulling the workflow out from existing systems so it can be examined on its own.

The objective is to understand where information moves, who makes decisions, and where delays or handoffs occur. To do that well, a collaborative team of functional experts, skilled engineers, and AI architects is needed at the table.

From there, leaders can ask a different question: If this workflow were designed today with AI-native capabilities in mind, what would it look like?

Only after the workflow has been reimagined does it get placed back into the business context. This exercise often surfaces opportunities to simplify work, reduce friction, and embed intelligence directly into execution. The process follows this pattern:

Choose a workflow:
Study it in isolation > Reimagine it > Reintegrate it

Then measure it, learn from it, and reuse what works in the next workflow.



Only after the workflow has been reimagined does it get placed back into the business context.

This first checkpoint is where many organizations discover that AI value is ultimately a workflow challenge. The unit of value is the workflow outcome, not the AI system. Once that foundation is in place, technology choices become significantly easier to evaluate.

From workflow improvement to workflow redesign

DIMENSION	PROCESS IMPROVEMENT	AI-NATIVE PROCESS REENGINEERING
Starting point	Existing workflow	Desired business outcome
Focus	Improving individual steps	Rethinking the flow of work end to end
Decision-making	Follows existing handoffs	Reconsiders where decisions occur
Knowledge flow	Information moves between separate teams and systems	Intelligence is available throughout the workflow
Outcome	Incremental gains	Workflow redesign and new operating capabilities



PODCAST
Voices from *The GuideLine*:
Moonshot first, then tools



GUEST
Ryan Cunningham
Corporate Vice President
Power Platform, Microsoft

"I think sometimes we get confused. We think, I need to go ship a chatbot, and that's the use case.

But all the interesting AI use cases don't start with technology at all. I met a customer in Australia last month who said, 'Here's my AI use case. Today, it takes me 36 hours to generate a quote. I want it to be two hours by June.'

That is an AI use case: 36 hours to two hours by June. That means I've front-loaded the ROI so I know what I'm willing to spend or risk to make that outcome happen.

I don't have to wait for a brainstorm and then some post-rationalization of AI value. I know the value upfront. Now I need to get really curious about how to make it happen."



[Listen to the full episode](#)



INSIGHT
Tokens: The economics of intelligence

A token is a basic unit of text an AI model reads and writes. Roughly speaking, one token represents about four characters, or three-quarters of a word in English. Every prompt, document, and response is measured and priced in tokens. More context can improve AI reasoning and decision-making, but it also increases cost.

That makes tokens an architectural decision, not just a technical one. Leaders should understand how workflows consume intelligence and how the organization pays for it. Consumption-based services typically price usage by tokens. Self-hosted or open-weight models shift that expense toward infrastructure, engineering, and operations rather than eliminating it. Most enterprises will use a portfolio approach: lower-cost or enterprise-controlled models for routine work, and frontier models where greater reasoning capability justifies the cost. The architectural task is to route work deliberately based on value, risk, latency, privacy, and required capability.

The architecture decision is about where your intelligence lives and who controls it.



CHECKPOINT 2

Build for intelligent execution

Do you know where your data lives?

With a workflow selected, the next checkpoint is designing the architecture to support it. The challenge is that AI is evolving faster than traditional technology cycles. Today’s frontier model may not be tomorrow’s. New providers will emerge. Costs will shift. Capabilities will improve. The organizations that create lasting value build architectures that can evolve without locking themselves into any one provider. Think of it this way: **Don’t pave the roads while the map is still being drawn.**

Fungibility and flexibility start with a clear understanding of where organizational intelligence sits. As workflows become more AI-enabled, you need to answer two critical questions:

- 1 | What data remains on-prem or within systems your organization directly controls?
- 2 | Which assets can be made available to external models and under what conditions?

There’s no single right answer, and the balance will differ by organization and by regulatory environment. What matters is the separation itself. Technologies such as data lakehouses, Model Context Protocol (MCP) servers, and shared context layers can help keep the enterprise intelligence layer distinct from the model layer. Your data, context, and accumulated judgment stay yours, while the models working on top of them can be swapped as better ones arrive. This is the digital brain taking shape. The architecture decision is about where your intelligence lives and who controls it.

This checkpoint is also the stage where the economics of AI come into focus. Every prompt, retrieval, and agent interaction consumes computational resources measured in tokens (*see Insight sidebar*). Understanding those costs early helps leaders make better decisions about governance, architecture, and scale.

This is assured architecture: composable, secure, governed, reusable, and designed to evolve without surrendering control of data, context, or institutional intelligence. It preserves model choice while keeping ownership of the knowledge and controls that make those models useful.

CHECKPOINT 3

Activate people + processes

An AI team can be a bottleneck

Once you’ve identified a workflow and the architecture is in place, the next challenge is building human systems that allow AI to scale. Value rarely comes from a dedicated AI team working in isolation. It emerges when the people who understand the workflow best are equipped with the tools and authority to improve it. Process owners, frontline practitioners, data leaders, technologists, and governance stakeholders all have a role in redesigning the workflow.

Many organizations also adopt a federated operating model. In practice, this means that **AI expertise is distributed throughout the business rather than concentrated in a single center of excellence.** Knowledge owners remain embedded in their functions while governance, platforms, and standards are managed centrally. This approach brings data and AI closer to day-to-day decision-making while maintaining consistency across the enterprise.

Workforce adoption follows a similar pattern. Instead of assigning participants, the strongest programs begin at the grassroots level, identifying people across the board who are eager to learn and experiment. Those early adopters become change agents by sharing what works with their peers. Structured training programs, partnerships with universities and platform providers, and hands-on project work help turn interest into capability.

The most effective programs connect that capability to consequential work. They give employees opportunities to apply AI to real business challenges, surface institutional knowledge trapped within teams and functions, and share what they learn back across the enterprise. Over time, those learning loops become feedback loops, helping the organization improve how the workflow gets done.



Value rarely comes from a dedicated AI team working in isolation.

PODCAST
Voices from *The GuideLine*:
A masterclass in people-first AI



GUEST
Jessica Baker
Chief Data Officer
Tennessee Valley Authority

“You must have the fundamental belief that AI doesn’t take people out of the equation. It’s *the* tool to empower people.

When we started our AI journey, we started with workforce, then moved into adoption, and then scaled value across TVA.

We started an AI change agent network. And it was just an open invitation—who is passionate about it, who wants to learn about it. We grew to 250 people across the company.

Then we developed a 12-week program. It starts with data and cyber responsibility. Then it gets more advanced: machine learning, Python coding. They sit in a classroom. They take exams. And they design a capstone AI solution for TVA by the time they graduate.

At the end they’re telling you, ‘I didn’t think this was possible. I didn’t think I could work this fast.’”



[Listen to the full episode](#)

Context makes amplifying expertise possible. It's the digital brain at work.



CHECKPOINT 4

Orchestrate agents

Don't forget your digital workers

As organizations redesign workflows around AI, they must decide how agents will contribute to execution. AI agents can retrieve information, make recommendations, and coordinate across systems. In many cases, they're beginning to function as members of an extended workforce.

This creates a new management challenge. Traditional software follows instructions. Employees exercise judgment. Agents sit somewhere in between. They're deployed like technology assets but operate with a degree of autonomy that requires supervision. The question becomes: How should work be divided between people and agents to create the greatest value?

The most successful organizations start at the workflow level. They break work into individual tasks and decide which responsibilities remain human-led, which can be delegated to agents, and where collaboration works best. **The goal is a clear division of labor: People focus on judgment, relationships, and improvement while agents handle process-driven tasks.** When done correctly, this approach amplifies expertise rather than replacing it.

This is assurance by design: authority, accountability, evidence, and escalation built into the workflow from the start. Governance defines agent identity, access, decision rights, consumption limits, human intervention, auditability, and performance monitoring as part of execution itself.

Context makes amplifying expertise possible. It's the digital brain at work (*see Insight sidebar*). Agents require access to business rules, permissions, and process knowledge. Without context, outputs will be generic. With it, agents can contribute meaningfully to the workflow.

But context isn't enough. Agents also need boundaries. If context helps an agent understand the work, governance determines the scope of that agent's authority. Leaders must determine what agents can access, where human review is required, how exceptions escalate, who's accountable for outcomes, and how performance will be monitored over time.

Good orchestration requires an operating model that enables people and agents to work together within clearly defined roles and guardrails.



FROM GUIDEHOUSE AI STUDIO Context and ontology

As AI agents become more autonomous, they need more access to data. They also need context to understand what that data means.

An ontology is like a shared dictionary for the organization. It captures the terms, metrics, relationships, and rules that help employees interpret information and make decisions. In an agentic environment, an ontology gives AI agents the same shared understanding of the business.

Like any business asset, an ontology requires active stewardship. Organizations should assign ownership to the teams closest to the work, establish clear definitions for key concepts and metrics, and regularly update the ontology as policies and priorities evolve. Ongoing governance helps ensure that agents use current business knowledge rather than outdated assumptions.

In regulated environments, the ontology also becomes a control surface, linking business meaning to policy, authority, evidence, and accountability.



[Read the full article](#)

Much like a pilot relies on instruments throughout a flight, leaders need ongoing visibility.

CHECKPOINT 5

Measure + scale value

Value comes through deliberate reinvestment

The final checkpoint brings the model full circle. When organizations select a workflow in the first checkpoint, they’re meant to define what success looks like upfront. The final step is proving that value was realized.

Many organizations measure AI activity instead of business impact. Adoption rates, prompts, and model usage can provide useful signals, but they rarely tell the full story. The most effective metrics are tied directly to workflow outcomes: speed, cost, quality, behavior change, business impact, assurance, and return on reinvestment (*see Insight sidebar*). Capacity is often the first value-add from AI: hours freed or backlogs cleared. Reinvestment is when leaders deliberately redirect that capacity into higher-value work, improved service, or greater mission throughput, then measure what changed.

It’s important to note that measurement shouldn’t end at the launch of the redesigned workflow. Establish a baseline before implementation, monitor results over time, and use what you learn to continuously improve the workflow. Much like a pilot relies on instruments throughout a flight, leaders need ongoing visibility into performance. This feedback loop strengthens both the process and the underlying AI capabilities.

Although the checkpoints in this model are presented sequentially, value creation is iterative. **As one workflow begins delivering measurable results, you can apply the same model to the next high-priority opportunity.** Over time, each successful implementation reinforces the next. As capabilities are folded in, intelligence strengthens, creating a scalable, repeatable cycle of learning, improvement, and value realization.



INSIGHT Metrics that matter

Every workflow redesign will have its own success metrics that are shaped by goals defined at the outset. The table below highlights some common categories and indicators used to track AI value.

MEASURE	WHAT TO LOOK FOR	EXAMPLE INDICATORS
Speed	Is work moving faster?	• End-to-end process time • Issue resolution time • Frequency of forecast updates
Cost	Are resources being used more efficiently?	• Operating cost • Technology spend (licenses, models, tokens) • Cost per unit of work (reviewing a claim or processing an invoice)
Quality	Are outcomes improving?	• Error rates • Accuracy (% of tasks that require no rework) • Compliance outcomes (audit findings or policy violations)
Behavior change	Are people working differently?	• Workflow usage (% of eligible work conducted through the new workflow) • Agent utilization (# of agent interactions) • Employee engagement (training completion rates, survey ratings)
Business impact	Is the workflow delivering its intended outcome?	• Revenue growth • Customer or constituent experience • Mission outcomes (outage response times or loan processing)
Assurance	Is the workflow operating within its intended boundaries?	• Policy and compliance exceptions • Human override and escalation rates • Time to detect and resolve failures
Return on reinvestment	Is newfound capacity being put to work?	• Hours redirected to higher-value work • New services or capabilities launched • Backlog reduction converted to output

VISUALIZATION

The path to the intelligent enterprise


Operationalizing intelligence requires a systematic approach to selecting the right work, redesigning how that work gets done, preparing people and agents, establishing governance, measuring outcomes, and reinvesting the capacity created.

The five checkpoints shown here provide a practical framework for doing just that.

At the center is the digital brain, the connected intelligence that learns and strengthens with every redesigned workflow. Surrounding the model is assurance, the guardrails that help the enterprise manage risk and maintain trust.

This isn't a one-time implementation path. Over time, the flywheel becomes a repeatable system for turning frontier capability into assured enterprise outcomes.

- 1 | **Define the work**
Start with the outcome, not the tool.
- 2 | **Build for intelligent execution**
Keep organizational intelligence connected, controlled, and ready to evolve.
- 3 | **Activate people + processes**
Embed AI skills and decision-making across the teams closest to the work.
- 4 | **Orchestrate agents**
Define how people and digital workers contribute.
- 5 | **Measure + scale value**
Show results and reinvest the capacity created.



3 Start here


Enterprise intelligence in action

Where value is showing up by industry

So, what does enterprise intelligence look like in practice? It depends entirely on the work. But it doesn't arrive all at once. It builds workflow by workflow as organizations redesign how they make decisions and take actions. For utilities, it may mean restoring power faster. For banks, it may mean moving a loan from application to decision in hours.

The pattern underneath stays consistent. An organization chooses a high-stakes workflow, reimagines it, redesigns it, and carries the learnings onto the next round.

The examples that follow spotlight where such gains are already taking hold today, illustrated by use cases with links to dive deeper.



Energy and utilities

For utilities, AI is helping improve operations, strengthen grid reliability, and enhance customer outcomes. Beyond automating tasks, the intelligence layer enables faster decisions, more accurate planning, and more effective responses to affordability pressures and burgeoning load growth from data centers, which will potentially consume as much as 15% of U.S. electricity by 2035, according to Guidehouse forecasts.

USE CASE

Utility customer experience

AI can help utilities simplify one of the industry’s most complex customer journeys: understanding and managing energy use. By combining data from meters, devices, and grid operations, a digital brain can deliver more personalized service and simpler billing.



[Read the full article](#)

Financial services

In financial services, many of the strongest opportunities exist in middle- and back-office functions where work is document-heavy, rules-driven, and full of exceptions. Loan processing, trade operations, AML investigations, fraud detection, regulatory reporting, and customer service are all areas where redesigned workflows can deliver measurable value.

USE CASE

Assured decision intelligence

Many financial institutions have digitized intake, but consequential decisions still depend on fragmented data, manual review, and difficult handoffs. AI-enabled decision intelligence can redesign lending, onboarding, fraud, and financial-crime workflows, bringing governed data, institutional context, human judgment, and traceable controls together. The result is faster decisions and fewer unnecessary reviews.



[Read the full article](#)





Government

For government agencies, the value unit is the mission outcome: benefits delivered faster, cases resolved sooner, services reaching the right people at the right time. Common starting points for the digital brain are high-volume workflows like claims processing, case management, grants administration, and constituent services.

USE CASE

AIOps for national security

Modern homeland security operations generate more system data than human teams can reasonably process. AI-powered operations, or AIOps, apply intelligence to detect anomalies, predict failures, and surface issues before missions are affected.



[Read the full article](#)

Healthcare

In healthcare, intelligence is showing up in administrative and middle-office work: prior authorization, scheduling, claims, and clinical documentation. With large volumes of information, clear rules, and frequent handoffs, these workflows are a good fit for AI-enabled redesign. This is especially true in revenue cycle management, where 78% of providers are using automation and AI to speed manual processes, according to a 2026 Guidehouse study.

USE CASE

Ambient listening

An AI agent listens to a patient’s visit, records what’s said, and surfaces the orders the clinician mentioned. Clinicians get back time as they shift from writing notes to analyzing notes. Governance is key. Consent, retention, and a path to flag errors must be integral to the design.



[Read the full article](#)



4 Clear the way

Misconceptions that can stall progress

Watch for potholes along the way

the five checkpoints show how to build enterprise intelligence. The industry examples show where it's already taking hold. The list that follows focuses on misconceptions about AI technology that have surfaced as tools become more widely adopted. They can slow progress once leaders begin the work.

These common myths may sound reasonable on the surface, which is why they persist. Some reinforce old ways of working. Others make pilots look like progress, and tools like strategy. Left unchallenged, they can disrupt momentum and keep enterprise intelligence from scaling.

6

myths that
can slow
the shift to
enterprise
intelligence

1

“AI can’t be trusted.”

Trust in AI is engineered, not assumed. Leading organizations implement risk-based governance plus workflow design controls and then evaluate AI on the results.

2

“Our vendor will deliver AI value.”

A single-vendor approach can limit options as AI evolves. A fungible architecture makes it possible to combine providers, adopt new capabilities, and avoid costly lock-in as needs change.

3

“Our AI portfolio is proof of progress.”

Hundreds of pilots may demonstrate activity, but activity isn’t value. Transforming one or two mission-critical workflows can unlock more impact than dozens of disconnected experiments.

4

“My IT team owns AI.”

AI is an operating model change, not a technology initiative. Your IT team can enable it, but the biggest gains come when business leaders, operators, and technologists move together.

5

“AI is mainly a workforce-reduction play.”

AI doesn’t replace expertise. It multiplies it. Paired with experienced professionals, AI compounds value. Strip the expertise out and it amplifies mistakes.

6

“Our existing governance can handle agentic AI.”

AI is dynamic. Its guardrails must be, too. Governance must address elements like agent access, consumption limits, and escalation paths—and evolve continuously as the work changes.



The organizations creating lasting value will be the ones that choose their direction with intention and begin the climb now.

Closing thoughts

Every organization starts from a different place. Some begin with operational efficiency. Others focus on customer experience or mission delivery. What matters most isn't where the journey starts but which path leaders choose.

The path of incremental improvement remains available. But only the second path to enterprise intelligence fundamentally changes how an organization operates. It builds the digital brain, redesigns mission-critical work around it, and reinvests the capacity it frees.

As AI becomes ever more deeply woven into the fabric of business and government, the window for waiting is beginning to close. The organizations creating lasting value will be the ones that choose their direction with intention and begin the climb now.

ABOUT THE GUIDE

Methodology and contributors

The Tech Guide reflects a practitioner perspective based on Guidehouse client work, interviews with business and technology leaders, applied research, external market evidence, and on-the-ground experience of professionals working across regulated commercial and public sector environments. It's intended as a practical decision guide rather than an empirical ranking or predictive technology forecast.

CONTRIBUTORS



Stuart Brown
Partner, Chief Transformation Officer



Ron White
Partner, Chief Information Officer



Greg Meyer
Partner, Global Technology Leader



Chris Scott
Partner, Chief Architect



Karen Odegaard
Partner, Global AI & Data Leader



Bob Audet
Partner, AI & Data, Defense & Security



Leigh Braunstein
Partner, Technology Leader, Health



Sourish Chaudhuri
Partner, Growth Leader, Financial Services



Grant Gengelbach
Partner, AI Studio Lead




Manish Kumar
AI Governance Program Office Lead



Molly Rupert-Sullivan
Director, AI Transformation



Leigh Sheldon
Partner, AI & Data, Communities, Energy & Infrastructure



Truc Taylor
Director, AI & Data, Health

Turn the Guide into results

The path to the intelligent enterprise begins with one consequential workflow and a clear definition of value. Guidehouse helps you identify that starting point, redesign the work, establish the required architecture, mobilize a cross-functional team, and scale proven results.

AI Studio helps turn promising ideas into informed action. AI engineers, data scientists, architects, and designers work alongside our industry and delivery teams to evaluate opportunities, rapidly prototype concepts, and address technical and governance considerations. This collaborative approach helps you reduce uncertainty, refine requirements, and identify the right path forward before committing to full-scale development.



[Explore the AI Studio](#)

THE TechGuide 2026

Guidehouse is a global professional services firm delivering advisory, technology, and managed services to the commercial and government sectors. With an integrated business technology approach, Guidehouse drives efficiency and resilience in the healthcare, financial services, energy, infrastructure, and national security markets.

[guidehouse.com](https://www.guidehouse.com)

© 2026 Guidehouse Inc. All rights reserved. This content is for general information purposes only and should not be used as a substitute for consultation with professional advisors.